



Dynamic Capacity Allocation Algorithms for iNET Link Manager

May 2014

**Tom Young
SET Executing Agent
412 TENG/ENI
(661) 277-1071
Email: tommy.young.1@us.af.mil**

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

**Test Resource Management Center (TRMC)
Test & Evaluation/ Science & Technology (T&E/S&T)
Spectrum Efficient Technology (SET)**

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 23-02-2015		2. REPORT TYPE Technical Paper		3. DATES COVERED (From - To) 8/09 -- 4/13	
4. TITLE AND SUBTITLE Dynamic Capacity Allocation Algorithms for iNET Link Manager				5a. CONTRACT NUMBER: W900KK-09-C-0021	
				5b. GRANT NUMBER: N/A	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Mariusz Fecko, Kirk Chang, Andrzej Cichocki, Heechang Kim, Shree Gadgil, Mohsen Sarraf, Melbourne Barton, Larry Wong, Sunil Samtani, Ray O'Connell, Bob O'Neil, Michael Rauf, Mark Radke, Tom Young, Thomas Grace				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) TT Government Solutions INC, 150 Mount Airy RD, Basking Ridge NJ 07920-2021 RoboComAI, 175 Tri County Parkway # 112, Cincinnati, OH 45246				8. PERFORMING ORGANIZATION REPORT NUMBER 412TW-PA-14256	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Test Resource Management Center Test and Evaluation/ Science and Technology 4800 Mark Center Drive, Suite 07J22 Alexandria, VA 22350				10. SPONSOR/MONITOR'S ACRONYM(S) N/A	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) SET 2015-0001	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release A: distribution is unlimited.					
13. SUPPLEMENTARY NOTES CA: Air Force Flight Test Center Edwards AFB CA CC: 012100					
14. ABSTRACT In an iNET telemetry network, Link Manager (LM) dynamically allocates capacity to radio links to achieve desired QoS guarantees. Under the T&E S&T iMANPOL program, we developed an enhanced capacity allocation algorithm that can better cope with severe congestion and misbehaving users and traffic flows. We compare the E-LM with the LM baseline algorithm (B-LM), which employs priority-weighted allocation. The B-LM is expected to perform well for the majority of traffic patterns, but does not prevent an ill-behaved traffic class from causing excessive latency on other radio links. The E-LM ensures that each class has a "guaranteed" portion of the total available bandwidth that is proportional to the weight of the class. If the traffic loading of a class is lower than its quota, the difference can be flexibly shared by other classes across multiple links. If the traffic loading of a class is higher than its quota, its demand may still be satisfied, provided that the capacity is not taken away from well-behaved traffic classes that stay below their quotas. The qualitative analysis shows the E-LM provides lower latencies for the well-behaved links in overloading conditions and increases the overall system throughput when the traffic is unbalanced. We conducted extensive experiments to confirm that analysis, with the E-LM reducing latency of well-behaved flows up to 90%, and increasing overall throughput up to 65% over the B-LM.					
15. SUBJECT TERMS Spectrum, Aeronautical telemetry, algorithm, bandwidth, attached sync marker (ASM), Integrated Enhanced Networked Telemetry (iNET), Link Manager (LM)					
16. SECURITY CLASSIFICATION OF: Unclassified			17. LIMITATION OF ABSTRACT None	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON 412 TENG/EN (Tech Pubs)
					19b. TELEPHONE NUMBER (include area code) 661-277-8615
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			

DYNAMIC CAPACITY ALLOCATION ALGORITHMS FOR INET LINK MANAGER

Mariusz Fecko¹, Kirk Chang¹, Andrzej Cichocki¹, Heechang Kim¹, Shree Gadgil¹, Mohsen Sarraf¹, Melbourne Barton¹, Larry Wong¹, Sunil Samtani¹, Ray O’Connell², Bob O’Neil², Michael Rauf², Mark Radke³, Tom Young⁴, Thomas Grace⁵

¹Applied Communication
Sciences, Basking Ridge,
NJ

²RoboComAI
Cincinnati, OH

³Tybrin
Corporation,
Edwards AFB, CA

⁴USAF AFMC,
Edwards AFB,
CA

⁵NAVAIR,
Patuxent River,
MD

ABSTRACT[†]

In an iNET telemetry network, Link Manager (LM) dynamically allocates capacity to radio links to achieve desired QoS guarantees. Under the T&E S&T iMANPOL program, we developed an enhanced capacity allocation algorithm that can better cope with severe congestion and misbehaving users and traffic flows. We compare the E-LM with the LM baseline algorithm (B-LM), which employs priority-weighted allocation. The B-LM is expected to perform well for the majority of traffic patterns, but does not prevent an ill-behaved traffic class from causing excessive latency on other radio links. The E-LM ensures that each class has a “guaranteed” portion of the total available bandwidth that is proportional to the weight of the class. If the traffic loading of a class is lower than its quota, the difference can be flexibly shared by other classes across multiple links. If the traffic loading of a class is higher than its quota, its demand may still be satisfied, provided that the capacity is not taken away from well-behaved traffic classes that stay below their quotas. The qualitative analysis shows the E-LM provides lower latencies for the well-behaved links in overloading conditions and increases the overall system throughput when the traffic is unbalanced. We conducted extensive experiments to confirm that analysis, with the E-LM reducing latency of well-behaved flows up to 90%, and increasing overall throughput up to 65% over the B-LM.

I. INTRODUCTION

In a multiple-access telemetry network such as the iNET [1] Radio Access Network (RAN), where an RF-link is shared across geographically dispersed nodes, allocating capacity to achieve QoS guarantee for multiple mission priority levels is a challenging task. The iNET network-based architecture provides this functionality through the Link Manager (LM) [2][3]. In the LM configuration, every QoS class is assigned a “class weight” based on DSCP, and every link is assigned “link priority weight” based on Mission Service Level Profile (MSLP) Weight/Priority. The LM instance at a ground node obtains the current per-mission/per-QoS class traffic demands and queue depths from airborne Test Articles (TAs) and ground network nodes (Figure 1). Using these inputs, LM acts as a TDMA controller to allocate slots by assigning RF channel capacity. It

[†] The authors would like to thank the Test Resource Management Center (TRMC) Test and Evaluation/ Science and Technology (T&E/S&T) Program for their support. This work was funded by the T&E/S&T Program through the U.S. Army Program Executive Office for Simulation, Training and Instrumentation (PEO STRI), Contract No. W900KK-09-C-0021. The Executing Agent and Program Manager work out of the AFTC. 412 TW-PA-14256

does so through generation of Transmission Opportunities (TxOps) messages that establish uplinks and downlinks and allocate transmission resources based on packet and mission priority.

The iNET Management and Operations with Policy Controls (iMANPOL) program developed several techniques to provide end-to-end QoS [4] for advanced telemetry networks. In particular, the iMANPOL capacity allocation algorithms are designed to help LM deal with difficult scenarios of severe congestion and ill-behaving users and traffic flows. The underlying premise is that protecting “well-behaved” links (i.e., the ones that do not overload the system at the expense of other links) and penalizing “ill-behaved” ones (i.e., the overloading links) is in line with demonstration/anticipated CONOPS for the test range networked telemetry system.

In the course of the project, these algorithms have been adapted for the LM architecture, implemented, and validated. The modeling and evaluation effort has confirmed the feasibility and value of the presented approach.

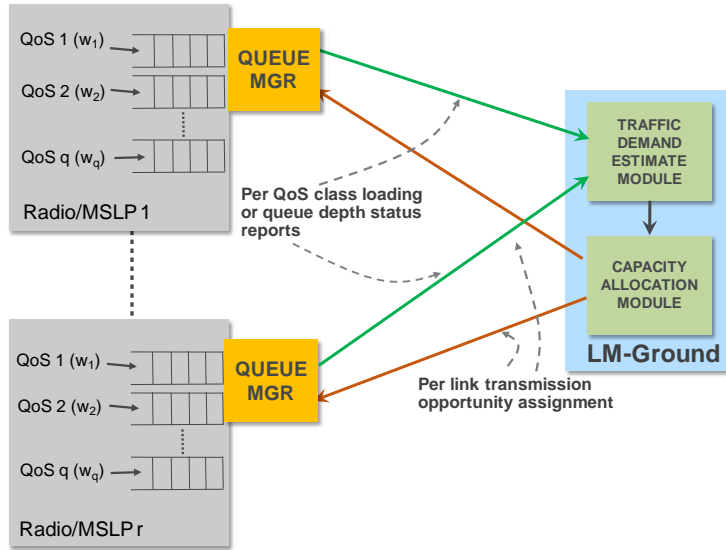


Figure 1: High-Level LM System Architecture

The rest of the paper is organized as follows. Section II describes and compares different capacity allocation algorithms. Performance evaluation study is presented in Section III. Section IV concludes the paper.

II. CAPACITY ALLOCATION

The goal of capacity allocation is to minimize latency when traffic that exceeds the current allocation is sent to the radio. If the allocation can be adjusted to traffic demand, it will reduce the possibility of data loss on radio queues due to overflow or timeout conditions.

Capacity allocation algorithms should follow sound Traffic Management principles to cope with three conditions: (1) Severe congestion, (2) Bursty traffic (VBR, On/Off), and (3) Ill-behaving users and traffic flows. (An ill-behaving flow is one that offers more traffic than indicated by its priority relative to other flows.)

A. Baseline Algorithm (B-LM)

The B-LM allocates on-demand capacity in proportion to the queue weight and the share of demand that the given queue contributes to its associated traffic class among all links. After minimum capacity is allocated to each traffic class (to guarantee basic fairness), the remaining on-demand capacity is iteratively allocated to each traffic class in proportion to traffic class capacity demand ratio:

$$\text{Demand ratio} = (\text{Link weight}) * \text{Class weight} * (\text{Class demand} / \text{Total class demand})$$

$$\text{Allocation} = (\text{Total capacity}) * \text{Demand ratio}$$

The B-LM algorithm is expected to perform well for the majority of traffic patterns.

B. Enhanced Algorithm (E-LM)

In the E-LM algorithm developed under the iMANPOL program, each class has a “guaranteed” portion of the total available bandwidth (called “quota”) that is proportional to the

weight of the class. Quota is computed dynamically as the fair share of currently available on-demand capacity.

When applied to high priority traffic, quota prevents possible starvation of low priority traffic, thereby ensuring fairness. The quota for high priority flows will be higher than for low priority flows proportional to the priority ratio, i.e., $quota_{high} / quota_{low} = weight_{high} / weight_{low}$. An ill-behaved class, i.e., the one exceeding its quota, may be degraded, but it should not adversely impact other classes. If the traffic loading of a class is lower than its quota, the difference is shared by other classes. If the traffic loading of a class is higher than its quota, its demand may still be satisfied, provided the loadings of some other traffic classes are less than their quotas.

These principles are applied and are the most effective in presence of multiple traffic classes and multiple links. The flow chart of the E-LM algorithm is shown in Figure 2.

In the first step, the algorithm allocates capacity equal to $\min(Quota, demand)$, where $Quota = (Total\ capacity) * (Link\ priority\ weight) * (Class\ weight) / (Total\ weight)$. The quota depends on configured link/class weights (fixed) and available capacity (varying). However, it does not depend on demand, which prevents greedy flows from capturing too much capacity in overload conditions. The latency experienced by non-overloading flows will thus be reduced thanks to the application of the quota.

In the second step, remaining capacity is distributed among queues according to demand and weight.

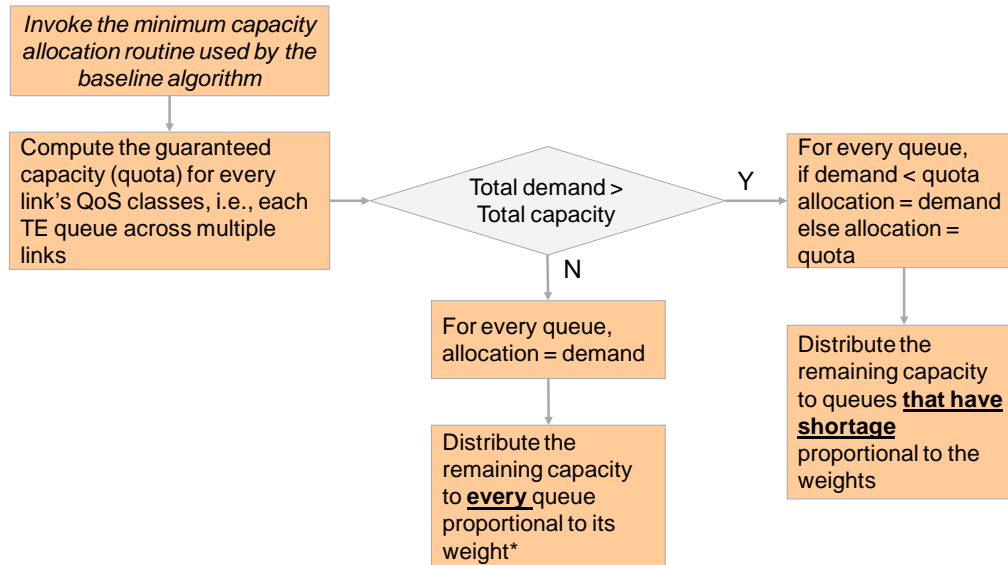


Figure 2: E-LM Algorithm

C. Qualitative Comparison

Compared with the LM baseline algorithm (B-LM), which is relatively simple to configure, the enhanced algorithm (E-LM) has more dependencies on various parameters in Layers 2 and 3. As to the performance, if the offered load does not exceed the total on-demand capacity, there should be no difference between the B-LM and the E-LM.

However, since the B-LM does not differentiate ill-behaved and well-behaved flows, it runs the risk of allocating capacity to well-behaved flows only after their queues experience excessive buildup. Consequently, an ill-behaved class may punish its own class in other missions/links, resulting in excessive latency. Additionally, capacity may be underutilized, reducing the overall

throughput of the system if loading of different traffic classes is “orthogonal” across links (i.e., each traffic class is on a separate link).

The E-LM should work well regardless of the overloading conditions as each traffic queue is guaranteed its quota. Only ill-behaved links and traffic classes with traffic loadings exceeding their quotas get penalized. The E-LM should protect well-behaved flows, and may satisfy demands of the ill-behaved flows if the loadings of some traffic classes are less than their quotas. Finally, the E-LM should significantly reduce latency of well-behaved queues because their quotas are guaranteed regardless of the presence of ill-behaved flows.

III. PERFORMANCE EVALUATION

We performed an extensive performance evaluation of both algorithms in two different platforms. A Linux testbed allowed us to use real traffic and queue implementation in a simplified framework, where the objective was to capture and verify main behavior aspects. However, the Linux testbed has the limitation of not modeling certain details such as:

- Timing of queue draining based on LM commands,
- Interplay between MAC and Traffic Engineering (TE) queues at the IP layer,
- Pre-built Code Blocks.

Hence, all test cases have been subsequently replicated in a higher fidelity OPNET LM environment.

A. Linux Testbed

To implement iNET Traffic Engineering (TE) Queues, the testbed shown in Figure 3 enhances the Hierarchical Token Bucket (HTB) queue provided in Linux kernel. To control the queue remotely, the LM sends UDP datagram (emulating TxOp) every 100ms to the VMs (emulating the TAs) containing capacity assignments. The HTB in the VM responds with its queue statistics (e.g., traffic loading and queue depth reports). The LM receives the reports, computes per-link capacity allocations using one of the two algorithms, and sends allocation results to the VMs in the next epoch. The HTB analyzes the datagram from the LM and drains packets from the HTB queue structure according to the per-link capacity allocations.




Figure 3: Linux Testbed

This approach provides better representation of actual traffic demands when the TE queues are saturated. It also accounts for traffic buffered at both TE and MAC queues. As shown in Figure 4, statistics are measured for incoming IP traffic (red arrow between “ip” and “net_intf”) at the input of the net_intf process model before IP packets are forwarded to each TE queue. We collected the following metrics:

We collected the following metrics:

Throughput: Average rate of data delivery, in megabits per second (Mbps).

Delay: Queuing delay in the Linux kernel, in milliseconds (msec).

Jitter: The average inter-packet arrival time measured at the destination in msec (as defined in IETF RFC 3550).

B. OPNET LM Testbed

In OPNET experiments, we used offered traffic estimates per TE queue instead of TE queue depth as the estimate of traffic demand.

Throughput: Application traffic received (per-queue) in bytes/sec. Transport Layer (UDP or TCP) traffic received (per-link), which will be forwarded to the application layer.

Latency (End-to-end Delay): Time taken for the packet to reach its destination, measured as the difference between the time a packet arrives at its destination and the creation time of the packet, measured in seconds (sec).

Queue Depth: Traffic Engineering (TE) queue depth at Test Articles.

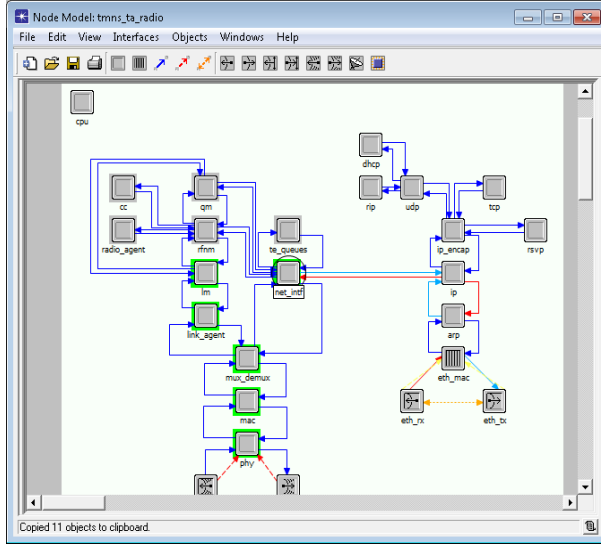


Figure 4: LM OPNET Model

LM Parameter	Value
TDMA Duration (Epoch)	100 msec
TDMA Guard Time	1 msec
Number of Active Links (down links)	3
Number of Queues (per link)	3
TE Queue Size	64 Kbytes
MAC Queue Size	32 Kbytes
Initial Working Capacity	8 Mbps
Minimum Capacity per link	493460 bps
Radio Channel Data Rate	34.36558 Mbps
X-Factor*: (PHY Layer Downlink Working capacity) / (IP Layer Received Traffic)	3.89

Figure 5: LM OPNET Configuration Parameters

To configure the E-LM algorithm in OPNET, we needed to appropriately map traffic load between IP and PHY layers. The E-LM uses the concept of ‘quota’ derived from the PHY layer “working capacity” resource. Hence it needs to know traffic demand at PHY layer, which it calculates by multiplying measured load (at the IP layer) by the so-called ‘X-factor’:

$$X = \{\text{all layers under IP layer overheads (MAC Frame Bytes, FEC Coding 2/3 Rate, Burst Sequence Header Bytes, ASM Bytes, etc)} + \text{IP packet size}\} / \text{IP packet size}$$
‘Effective PHY load’ = ‘X’ times ‘IP layer load’

The result is a traffic demand viewed in the PHY layer. Inside the E-LM algorithm, this PHY layer traffic demand is compared with the ‘quota’. The value of the ‘X-factor’ was derived both analytically and experimentally and added to the LM parameter set (Figure 5).

C. Test Cases and Performance Results

We defined a number of test cases summarized in Figure 6: Test Cases 1–4 are tailored to create overload conditions in selected queues, mixing CBR, VBR, and bursty ON-OFF traffic; Test Case 5 is tailored to show handling of highly unbalanced traffic loading both within and across links; and Test Case 6 replaces “over-loading UDP traffic” with TCP flow in one queue.

Consider Test Case 1 (Figure 7), which uses customized CBR test flows with periodic packet arrival and fixed-size packets of 8 KB. The total on-demand capacity across all links and queues is 8Mbps, whereas the total offered load is 12.8Mbps. Links 2 and 3 are below their load quota and Link 1 (Queue 1) exceeds its load quota.

Test-case	Traffic Loading	Traffic Pattern*	Overloaded Link(s)
0	Customized test flows	CBR	None (sanity check)
1	Customized test flows	CBR	Hi-priority link 1
1.1	(1) with addition of ON-OFF traffic	CBR	Hi-priority link 1
2	Customized test flows	CBR	Lo-priority links 2 and 3
2.1	(2) with addition of ON-OFF traffic	CBR	Lo-priority links 2 and 3
3	Customized test flows	VBR	Hi-priority link1
3.1	(3) with addition of ON-OFF traffic	VBR	Hi-priority link1
4	iNET flows	VBR and CBR	Lo-priority links 2 and 3
5	Customized test flows with orthogonal loading**	CBR	None
6	Customized flows with one TCP flow	CBR and TCP	None (except a TCP flow)

Figure 6: Test Cases Synopsis




Figure 7: Configuration of Test Case 1 – High-loading flows shown in red

As depicted in Figure 8 and Figure 9, the Enhanced-LM and Baseline-LM show similar throughput performance, with the E-LM showing small increase for the low-loading links (7% overall, more than 16% for the impacted traffic class in Queue 1) and 17% decrease for the high-loading link (limited to the offending class in Queue 1).

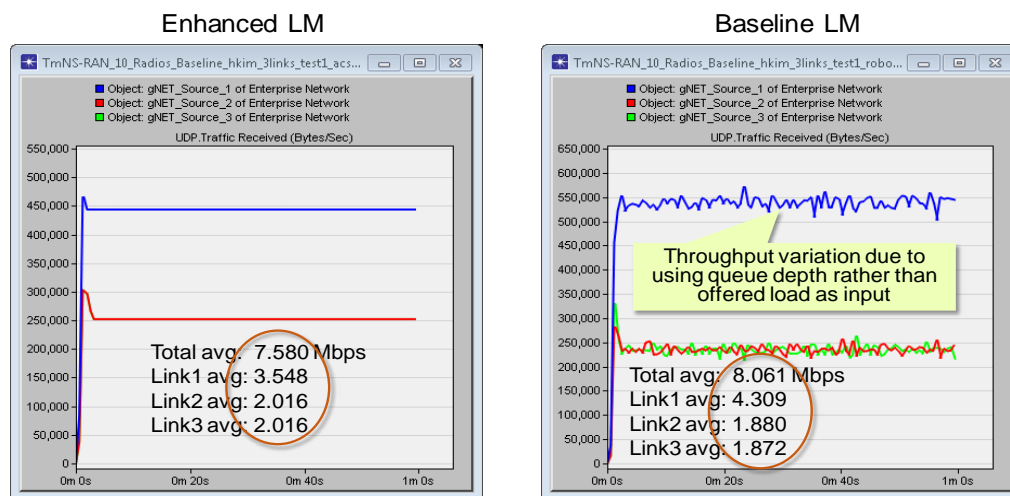


Figure 8: Test-Case 1 Results – Per-Link UDP Throughput Performance (OPNET)

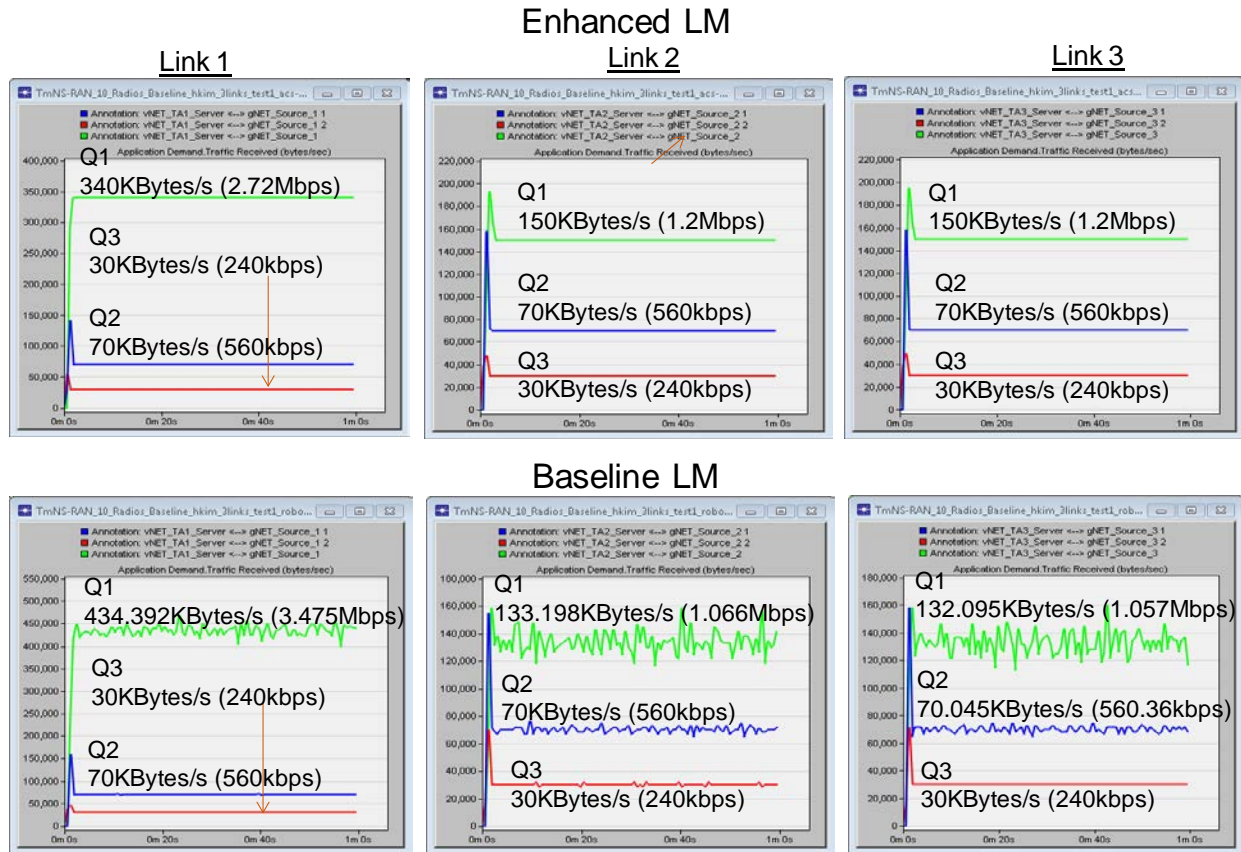


Figure 9: Test-Case 1 Results – Per-Queue Application Demand Throughput Performance (OPNET)

Figure 10 shows that the enhanced-LM significantly (by 76%) reduces average latency on well-behaved/low-loading links (Links 2 and 3), while latency on the overloaded Link 1 is increased by 27%. The observed latency improvements for the E-LM are achieved by the significantly reduced queue depths (Figure 11) on well-behaved/low-loading links.

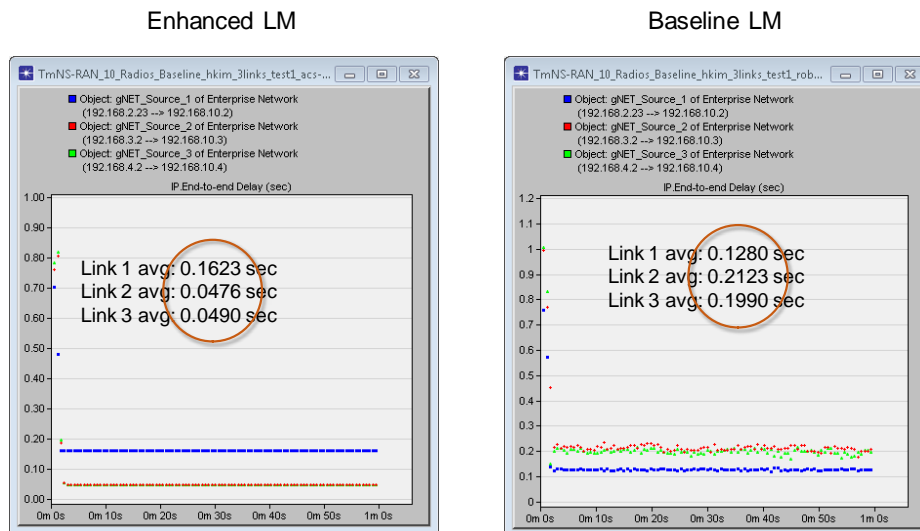


Figure 10: Test-Case 1 Results – Latency Performance (OPNET)




Figure 11: Test-Case 1 Results – TE Queue Depth Performance (OPNET)

The experimental results confirmed our qualitative analysis. When the high priority link (Link 1) was overloaded, the E-LM reduced overall latency (by >38% on average). For traffic on the low priority, well-behaved links (Link 2 and 3), the latency decrease is more significant (>76% on average), while the high priority, ill-behaved traffic's latency was increased (~27%). Latency decrease is achieved with small degradation to the throughput performance on overloaded link, and with a small increase of throughput performance on well-behaved links.

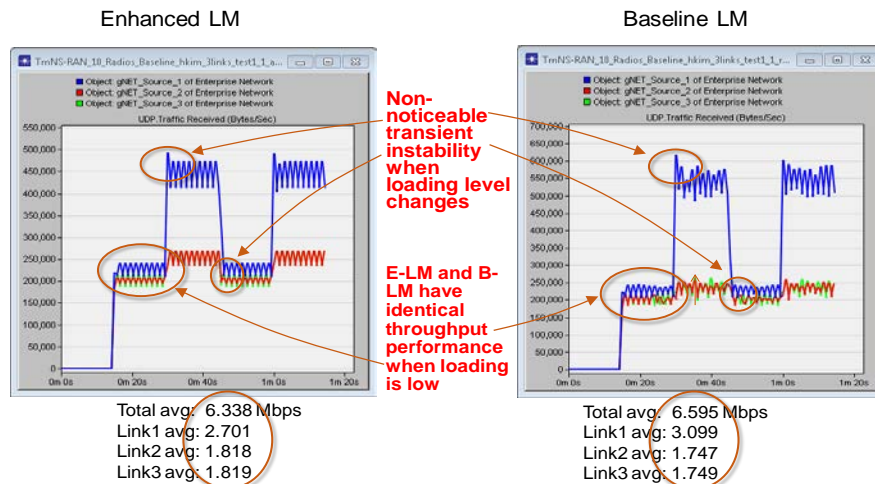


Figure 12: Test-Case 1.1 Results – Per-Link UDP Throughput Performance (OPNET)

Test Case 1.1 is an extension of Test Case 1 with an ON-OFF traffic model. The E-LM and the B-LM show similar throughput performance, with the E-LM producing small increase for the low-loading links and small decrease (13%) for the high-loading link (Figure 12). The E-LM

also significantly reduces average latency on well-behaved/low-loading links (Link-2 and 3) by 60%, while latency on overloaded Link 1 increased by 24% (Figure 13). In both test cases, the latency decrease is achieved without degrading throughput or jitter performance (Figure 14).

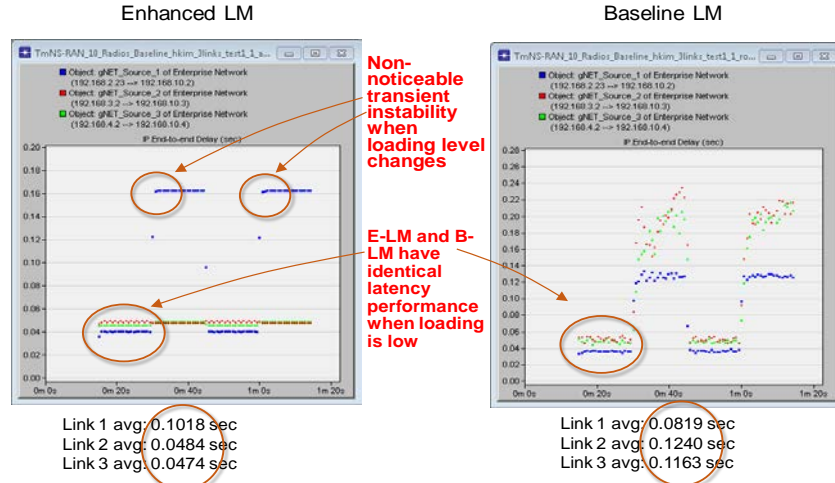


Figure 13: Test-Case 1.1 Results – Latency Performance (OPNET)

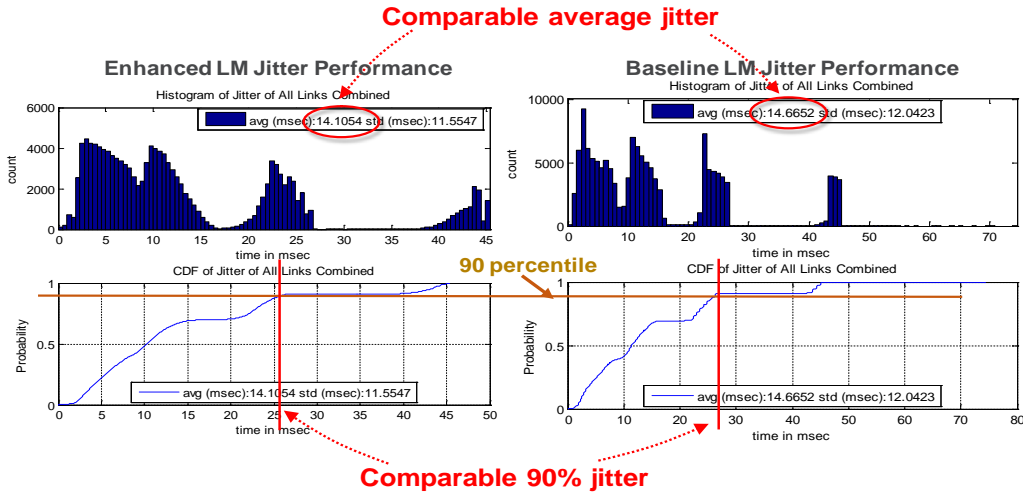


Figure 14: Test-Case 1 Results – Jitter Performance (Linux)

Test Case 5 uses highly unbalanced, but not overloaded traffic with only one queue active per link, i.e., “orthogonal” loading: Link 1, Queue 1 (0.4 Mbps); Link 2, Queue 2 (3.6 Mbps); and Link 3, Queue 3 (3.6 Mbps). The total offered load is 7.6 Mbps vs. 8 Mbps of the available on-demand capacity. Link 2 and 3 are over their load quota and Link 1 is below the load quota. As shown in Figure 15, the E-LM significantly increases overall throughput (by >65%), with the E-LM and B-LM showing similar throughput performance for the low-loading Link 1.

IV. CONCLUSION

The Enhanced-LM clearly performs better for unbalanced, overloaded systems. By utilizing the concept of allocation “quota,” this algorithm protects well behaved traffic flows. Both the qualitative analysis and the experimental results (in Linux testbed as well as OPNET environment) showed that E-LM algorithm provides better protection for the low loading links in overloading conditions than Baseline-LM. The E-LM shows significantly better latency

(improvements vary from 6% to over 60% in our experiments) for low-loading traffic without degrading its throughput and jitter. For traffic in well-behaved links, the latency decrease is even more significant (by 25%–90% on average). These results hold for both CBR and VBR traffic with multiple ON-OFF traffic scenarios (i.e., periodic bursting). E-LM also significantly increases throughput (>65%) for an extremely unbalanced traffic loading and does not cause transient instability when traffic loading level changes. The experiments also demonstrated that having even a single TCP flow can cause overload conditions. In this case, the E-LM better protects the low-loading UDP flows, similar to those non-TCP test-cases.

Both baseline and enhanced algorithms can coexist in the Link Manager, either statically configured or dynamically switched depending on traffic conditions. In the latter case, additional logic is needed to detect when the system is under stress due to (1) severe overload or (2) highly unbalanced traffic patterns. If such a situation occurs, the E-LM is activated. When traffic volume/patterns return to a normal operational regime, the B-LM is re-activated.

At the conclusion of the iMANPOL program, the E-LM has been integrated in the LM OPNET model that can be used to generate the operational code for a target deployment platform.ⁱ

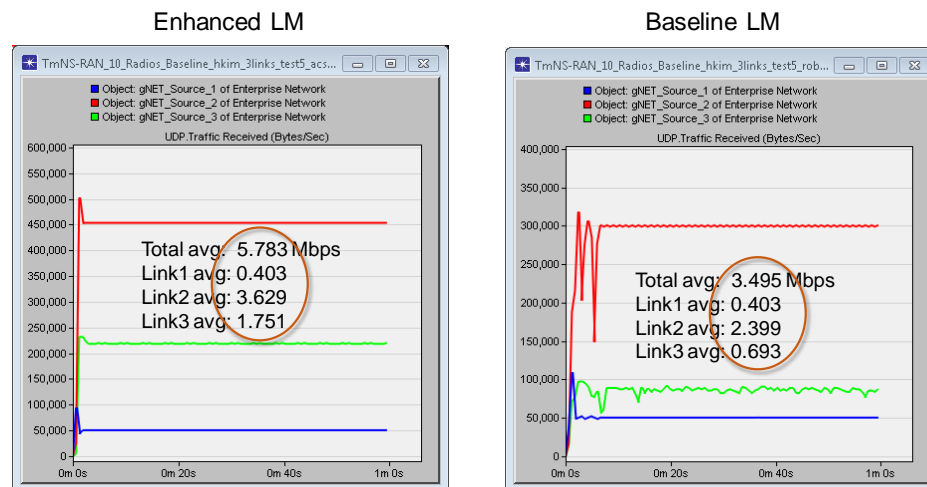


Figure 15: Test Case 5 Results: Per-Link UDP Throughput Performance

REFERENCES

- [1] B. Abbott, M. Araujo, M. Moodie, T. Newton, T. Grace. iNET System Design Concepts. In *Proc. ITC/USA International Telemetering Conference (ITC)*, 2011.
- [2] R. O'Connell, L. Webster, J. Kaba. Telemetry Network System (TmNS) RF Link Management Quality of Service. In *Proc. ITC/USA International Telemetering Conference (ITC)*, 2012.
- [3] R. O'Connell. Telemetry Network System (TmNS) Link Management Algorithm Verification. In *Proc. ITC/USA International Telemetering Conference (ITC)*, 2013.
- [4] A. Cichocki, M.A. Fecko, J. Unger, S. Samtani, L. Wong, A. Kolarov, M. Radke, T. Young. Dynamic end-to-end QoS management for advanced RF telemetry networks. In *Proc. ITC/USA International Telemetering Conference (ITC)*, 2011.

ⁱ Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the TRMC and T&E/S&T Program and/or the PEO STRI.